

**Technical Report
1205**

A Simple Probabilistic Combat Model

S.D. Weiner

DD Month YYYY

Lincoln Laboratory
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LEXINGTON, MASSACHUSETTS



This material is based on work supported under
Air Force Contract No. FA8721-05-C-0002 and/or
FA8702-15-D-0001.

This report is the result of studies performed at Lincoln Laboratory, a federally funded research and development center operated by Massachusetts Institute of Technology. This material is based on work supported by the Air Force Life Cycle Management Center/PZM under Air Force Contract No. FA8721-05-C-0002 and/or FA8702-15-D-0001. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the U.S. Air Force.

© 2016 MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Delivered to the U.S. Government with Unlimited Rights, as defined in DFARS Part 252.227-7013 or 7014 (Feb 2014). Notwithstanding any copyright notice, U.S. Government rights in this work are defined by DFARS 252.227-7013 or DFARS 252.227-7014 as detailed above. Use of this work other than as specifically authorized by the U.S. Government may violate any copyrights that exist in this work.

**Massachusetts Institute of Technology
Lincoln Laboratory**

A Simple Probabilistic Combat Model

*S.D. Weiner
Group 31*

Technical Report 1205

DD Month YYYY

Lexington

Massachusetts

This page intentionally left blank.

ABSTRACT

The Lanchester combat model has been used extensively to model the results of combat between two sides with (potentially) different quantities and qualities of forces. The model is deterministic; assuming that red attrition is proportional to blue force size and quality, and similarly for blue attrition. In the model described here, all attrition is modeled probabilistically and it is possible (although unlikely) for the weaker side to be successful. The model consists of a number of discrete waves in which red and blue forces attack each other. Since each attack has a probabilistic outcome, after each wave there will be a probability distribution of red and blue survivors. This distribution serves as the input to the next wave. For each wave, the model determines the discrete output probability distribution for each possible input of red and blue weapons. This discrete output probability is then convolved with the input probability to get the resulting overall output probability of red and blue survivors. This process is repeated for as many waves as needed to determine the probability that red or blue will win the battle. Results are shown for a variety of cases.

This page intentionally left blank.

TABLE OF CONTENTS

	Page No.
Abstract	iii
List of Illustrations	vii
1. INTRODUCTION	1
2. DETERMINISTIC LANCHESTER MODEL	3
3. PROBABILISTIC ISSUES	5
4. RESULTS AND TRADE-OFFS	9
4.1 Effect of More Unbalanced Engagements	10
4.2 Effect of Larger Forces on Both Sides	11
4.3 Effect of Increased or Decreased Lethality	11
4.4 Effect of Assignment Approach	13
5. MODEL FOR RED AND BLUE WIN PROBABILITY	15
5.1 Modeling the Probabilistic Results	17
6. SUMMARY	23

This page intentionally left blank.

LIST OF ILLUSTRATIONS

Figure No.		Page No.
1	Sample results for deterministic model.	3
2	Probabilistic model.	6
3	Possible allocations and their probabilities (10 blue targets).	7
4	Possible allocations and their probabilities (4 blue targets).	7
5	Sample probabilistic case.	9
6	Results for stronger red force.	10
7	Results for larger force levels.	11
8	Results for reduced lethality.	12
9	Results for increased lethality.	12
10	Both sides use random assignment.	13
11	Red random, blue coordinated.	14
12	Red coordinated, blue random.	14
13	Contours of red win probability.	15
14	Red win probability for different weapon allocations.	16
15	Red win probability for equal lethalties.	16
16	Red win probability vs. effectiveness ratio.	17
17	Red win probability for different force levels.	18
18	Approximate model.	19
19	Overall approximate model.	19
20	Comparison approximate model.	20
21	Approximate model for large force levels.	21

This page intentionally left blank.

1. INTRODUCTION

The Lanchester combat model¹ is a simple way to assess the effects of quantity and quality differences in opposing forces. In its simplest version, it is a set of coupled differential equations in which the rate of attrition of blue is proportional to the size of the red force, and the rate of attrition of red is proportional to the size of the blue force. The coefficients of proportionality depend on the qualities of the respective forces. In this form, the Lanchester model is deterministic. Given initial force levels and qualities, the balance of forces will evolve with time and, if neither side breaks off the combat, will result in the complete elimination of one (or both) of the forces at a finite time and with a deterministic force level for the surviving side.

There are a number of aspects of this model that are oversimplified, so probabilistic effects are introduced make it more realistic. These effects are most important when the forces' levels are modest, and effects of quantization (only integer number of forces can occur) are significant. This is the case for small numbers of combatants, as in aerial dogfights or tank battles. Probabilistic effects are introduced as simply as possible.

The rest of this report is organized as follows. Next, the deterministic Lanchester model is reviewed in order to develop the format used for presenting results. After that, the probabilistic model is introduced where, rather than a steady decline in each side's force level, there are a series of "waves" in which each blue weapon fires at a red weapon and each red weapon fires at a blue weapon. The outcome of each firing is a kill with a certain probability, and a miss with one minus the kill probability. Two firing doctrines are considered: random, where each target is selected at random; and coordinated, where all the firings by a side are assigned to make the attack as uniform as possible. This coordinated firing avoids problems with some targets not being attacked while others are attacked multiple times. This leads to an overall model in which the probability of red and blue survivors will evolve over time. For each combat wave, an input probability distribution is taken and transformed into an output probability distribution. A general format is presented to show results on how the probability distributions evolve, and also probabilities that the current status is more favorable to red or to blue. Finally, we show results for a variety of parameter sets and compare the results with those of the deterministic model.

¹. A good introduction to the model(s) is found in J.S. Przemieniecki's *Mathematical Methods in Defense Analysis*, 3rd ed. Reston, VA, USA: AIAA, 2000; and also in many Wikipedia articles.

This page intentionally left blank

2. DETERMINISTIC LANCHESTER MODEL

Let B be the level of blue force with an initial value of B_0 . Similarly, let R be the level of red force with an initial value of R_0 . We take the rate of change of B to be proportional to R with a coefficient Pkr , with Pkr being related to the effective kill probability. We take the rate of change of R to be proportional to B with a coefficient Pkb . This results in the coupled differential equations

$$\begin{aligned} dB/dt &= -Pkr R \\ dR/dt &= -Pkb B. \end{aligned} \quad (1)$$

Blue's losses are proportional to red's numbers and vice versa. Eq. (1) can be integrated to

$$Pkb(B_0^2 - B^2) = Pkr(R_0^2 - R^2). \quad (2)$$

The battle will be a tie if $Pkb B_0^2 = Pkr R_0^2$. If blue is stronger ($Pkb B_0^2 > Pkr R_0^2$), blue will have $\sqrt{B_0^2 - (Pkr/Pkb)R_0^2}$ weapons left when R is reduced to zero, should red not break off before this point. Figure 1 shows results using Eq. (2) for three different values of Pkr/Pkb . The curves show the evolution of the force levels (from upper right to lower left) for different initial values of the blue force. The shaded regions show which of the combatants is favored in the battle. Note that if we start in the blue region, we end in the blue region and vice versa. As the battle progresses, the outcome departs more and more from the break even line. As the kill ratio increases (red becomes more effective), a much larger blue force is needed to overcome red's quality advantage. A similar figure format is used for the probabilistic model, but there is the possibility of crossing the break-even line.




Figure 1. Sample results for deterministic model.

This page intentionally left blank

3. PROBABILISTIC ISSUES

The deterministic model assumes that the attrition of the forces occurs continuously and at a rate determined by the level of the opposing force. In reality, destruction of weapons is a discrete process that may or may not occur at a fixed rate. For small force levels, the difference between one weapon surviving or being destroyed can have a significant effect on future combat. In the probabilistic model, the probability distributions for the number of red and blue survivors is calculated as the battle progresses. This will be analogous to the deterministic Lanchester model, in that each red weapon will target a given blue weapon and vice versa. The difference will be that the target will be destroyed with a given probability Pk and will survive with probability $1-Pk$. In the deterministic model, the number of targets will be decreased by a given amount resulting in fractional numbers of targets surviving at any time.

Two types of targeting of red and blue weapons are considered. We start with a given number of red and blue surviving weapons. First, let's consider R red weapons attacking B blue weapons. If the attack is coordinated, red will try to attack blue as uniformly as possible to avoid overkill, leaving some targets not attacked. If $R < B$, then R targets will be attacked with one shot each and $B - R$ targets will not be attacked. If $R > B$, then we define $nlo = \text{int}(R/B)$ and $nhi = nlo + 1$. In the most uniform attack, $bhi = R - B \cdot nlo$ targets will get nhi shots, and $blo = B - bhi$ targets will get nlo shots. The survival probability for each bhi target will be $1 - (1 - Pkr)^{nhi}$ and the survival probability for each blo target will be $1 - (1 - Pkr)^{nlo}$. With these survival probabilities, the distribution of survivors in the blo class and in the bhi class can be calculated. These probabilities are given by the binomial distribution. These distributions can then be combined to get the overall distribution of blue survivors. We do the same thing for red survivors to get a combined distribution of red and blue survivors after an attack wave. It is easy to calculate the probability distribution of survivors when starting with fixed values of R and B . After the first attack wave, there will be probability distributions of R and B . Subsequent waves can be calculated by the distribution of R and B for each possible starting R and B values, and multiplying that resulting distribution by the initial distribution of R and B . In this way, we take an input probability distribution and transform it into an output probability distribution. This process can continue for as many waves as required. Figure 2 illustrates this process schematically.

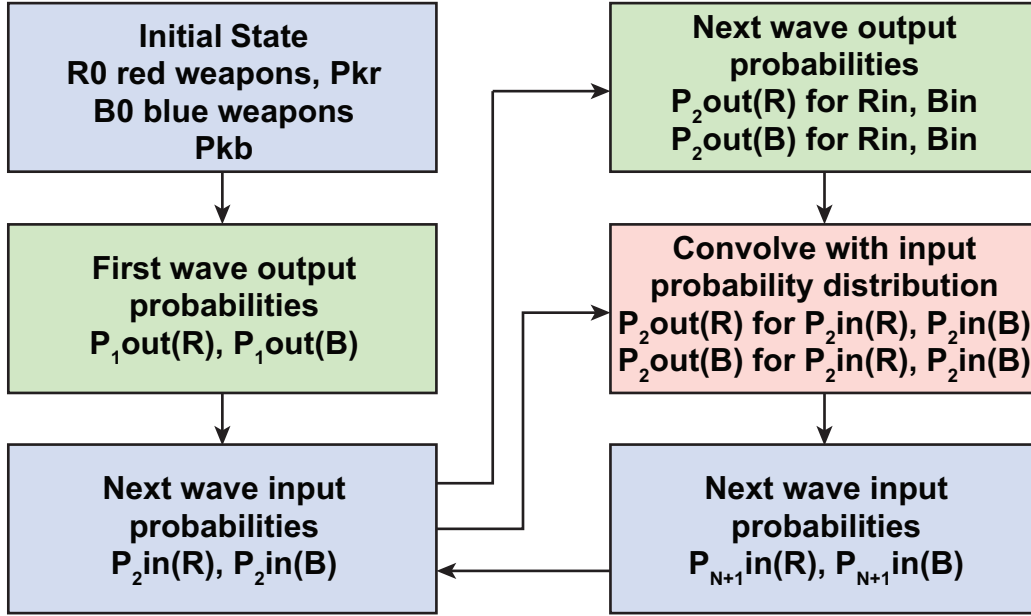


Figure 2. Probabilistic model.

The green boxes take a fixed number of input weapons on both sides with corresponding kill probabilities and calculate the distribution of surviving red and blue weapons. The pink box convolves these output distributions for specific input force levels with the actual input distribution of force levels, to get the resulting output distribution. The green and pink boxes together transform an input distribution into an output distribution. The blue boxes in Figure 2 just keep track of the bookkeeping of input and output distributions.

So far, the coordinated allocation of weapons in an attack wave have been considered. There can also be random allocation; where each red weapon picks a blue target independently from what the other red weapons do (and vice versa). In addition, there can be cases where both sides use coordinated allocation, where both sides use random allocation, and cases where one side uses coordinated allocation and the other side uses random allocation. Next, we will describe our random allocation case model.

For the random case, assume R red weapons are allocated to B blue weapons randomly. We are interested in the distribution of weapons assigned to the various blue targets. The constraint to use a total of R weapons makes some allocations more likely than others. Each possible allocation will have a given probability and will result in a given probability distribution of survivors. The survivor distributions can be combined with the probabilities of the allocations to get the resulting output probability distribution. Again, we have to convolve this output distribution for fixed input weapon levels with the actual input distribution to get the correct output distribution for the wave. Only the green boxes in Figure 2 will differ between the coordinated and random allocations.

To get the probabilities and kill probabilities for random allocation, it is useful to look at all the possible assignments of R red weapons to B blue targets. We will do this by ordering the blue targets by how many weapons are assigned to them. Figure 3 shows all the assignments and their probabilities for a sample case of 6 red weapons assigned to 10 blue targets

Allocation	Probability
6 0 0 0 0 0 0 0 0 0	0.0000
5 1 0 0 0 0 0 0 0 0	0.0005
4 2 0 0 0 0 0 0 0 0	0.0014
3 3 0 0 0 0 0 0 0 0	0.0009
4 1 1 0 0 0 0 0 0 0	0.0108
3 2 1 0 0 0 0 0 0 0	0.0432
2 2 2 0 0 0 0 0 0 0	0.0108
3 1 1 1 0 0 0 0 0 0	0.1008
2 2 1 1 0 0 0 0 0 0	0.2268
2 1 1 1 1 0 0 0 0 0	0.4536
1 1 1 1 1 1 0 0 0 0	0.1512

Figure 3. Possible allocations and their probabilities (10 blue targets).

It is very unlikely that all the red weapons will be assigned to just 1 or 2 of the blue targets, but it is more likely that 1 or 2 blue targets will receive multiple attacks, than that the attack will be uniform over the blue targets. Figure 4 shows a case with more reds than blues, where 6 red weapons attack 4 blue targets.

Allocation	Probability
6 0 0 0	0.0010
5 1 0 0	0.0176
4 2 0 0	0.0439
4 1 1 0	0.0879
3 3 0 0	0.0293
3 2 1 0	0.3516
3 1 1 1	0.1172
2 2 2 0	0.0879
2 2 1 1	0.2637

Figure 4. Possible allocations and their probabilities (4 blue targets).

Again, the most likely assignments are not the most uniform ones. For each allocation, there is a probability distribution for the number of surviving targets. By convolving these distributions with the probability of the allocation, the output distribution of survivors for that input number of attacking weapons can be determined. Finally, convolving that distribution with the input distribution of attackers, we can get the output distribution for a given input distribution. We will look at the differences in results for coordinated and random attack allocations. In both cases, the block diagram in Figure 2 is used to get the survival distribution after each attack wave.

4. RESULTS AND TRADE-OFFS

*The race is not always to the swift nor the fight to the strong,
but that's the way to bet.*

Hugh Keough

To illustrate some of the issues for probabilistic combat models, we will present results in a format analogous to that for the deterministic model. Figure 5 shows a sample result for an initial force of 10 reds and 5 blues. We assume the blue force is more effective (higher Pk) than the red force and look at how the survivor distributions evolve over the different waves of the battle.

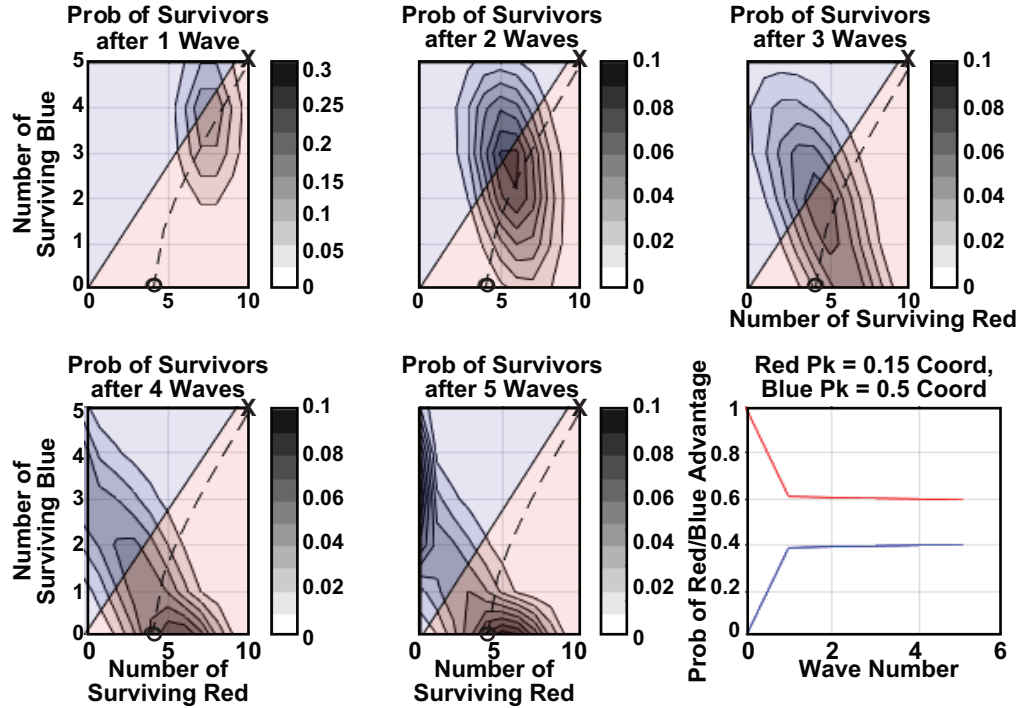


Figure 5. Sample probabilistic case.

The grid is the same as for the deterministic model, and the dashed line corresponds to the deterministic result for the given force levels and qualities. The red and blue shaded regions correspond to the regions where red or blue would be expected to win. The new information is the probability contours showing the probability of a given number of red and blue survivors after a given number of attack waves. Here we see quite a difference from the deterministic results. There is a significant spread in the number of survivors for both red and blue. The probability distributions cross over the break-even line indicating that

both red and blue have a significant chance of winning. By integrating the probability density over the blue and red shaded regions, we can get an estimate of the probability that blue or red will win the battle. These integrals are plotted as a function of the number of waves in the lower right plot. Since we start at a point in the space, there is a clear winner prior to the first wave. However, the uncertainty introduced by the outcome of the first wave leads to a close to 50–50 outcome. This is not surprising since the initial condition is very close to the break even line. What is more interesting is that the probability density tends to concentrate at either a clear red win or a clear blue win as the number of waves increases. This reflects the fact that, if a random fluctuation tends to favor red or blue, subsequent attack waves tend to amplify this fluctuation. In the following, we will vary a number of parameters to get a feeling for how the probability densities and win probabilities behave. We will look at 4 major topics: how things change as we get farther from the break even line, how things change as we increase the force levels, how things change as we increase the lethality levels, and how things change as we go from coordinated assignments to random assignments.

4.1 EFFECT OF MORE UNBALANCED ENGAGEMENTS

In Figure 6, we keep the same force levels but increase the relative quality of the red forces. We see that the probability density tends to follow the deterministic model (dashed line) fairly closely and red has a nearly sure win.

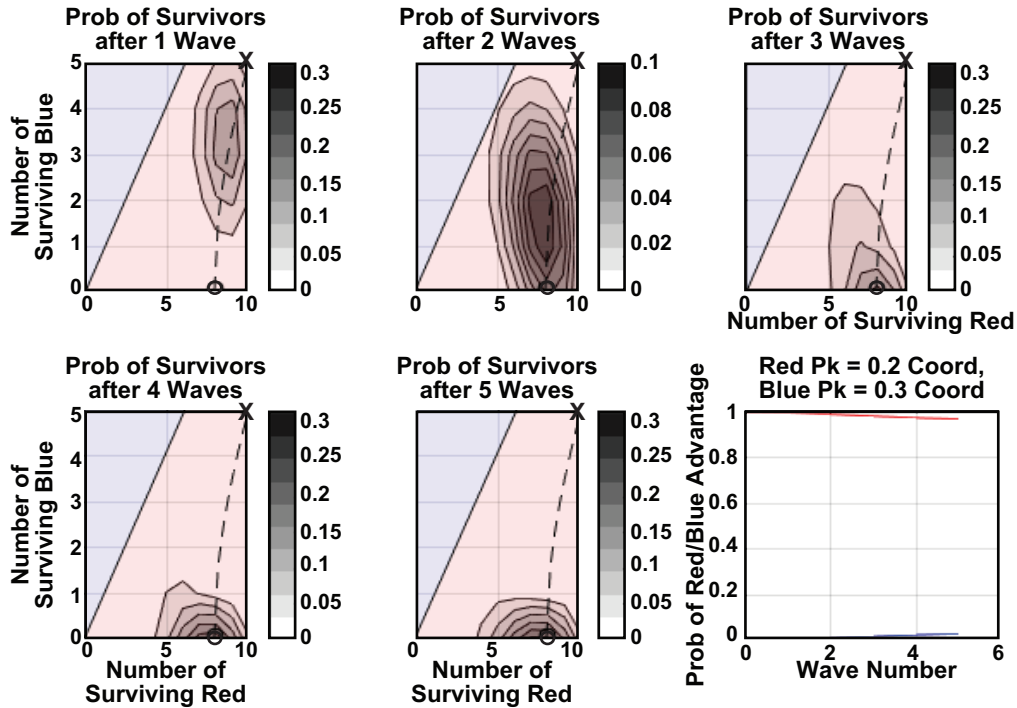


Figure 6. Results for stronger red force.

4.2 EFFECT OF LARGER FORCES ON BOTH SIDES

It is possible that the spread in results is a reflection of the small sizes of the red and blue forces. To get a feeling for this effect, Figure 7 shows the results for both forces being doubled from those in Figure 5. We see that there isn't much difference in the probability density or in the overall result. Later, we will look at even longer forces and will see a more significant effect.

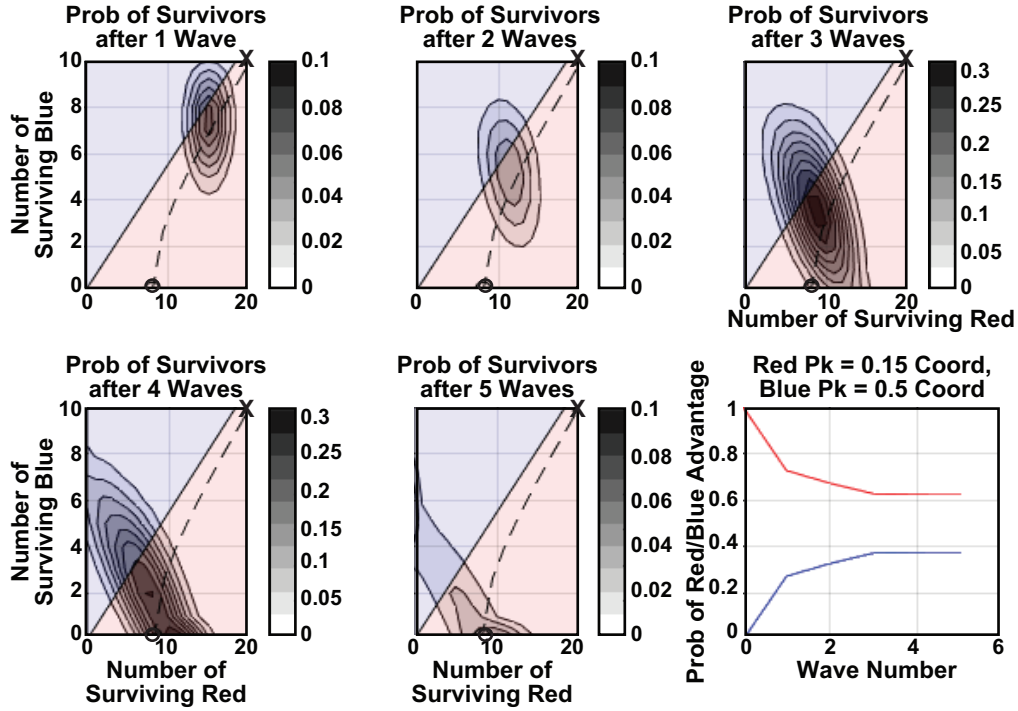


Figure 7. Results for larger force levels.

4.3 EFFECT OF INCREASED OR DECREASED LETHALITY

In Figure 8, we cut the lethality of both red and blue by a factor of 2. We see that the general behavior of the probability density and the win probabilities is similar to that in Figure 5 but the evolution occurs more slowly. The results after 4 waves with lower lethality are similar to those after 2 waves with higher lethality.

In Figure 9, we increase the lethality of both sides by 50%. Again, the probability density and win probabilities are similar to those in Figure 5, but the evolution occurs more rapidly.

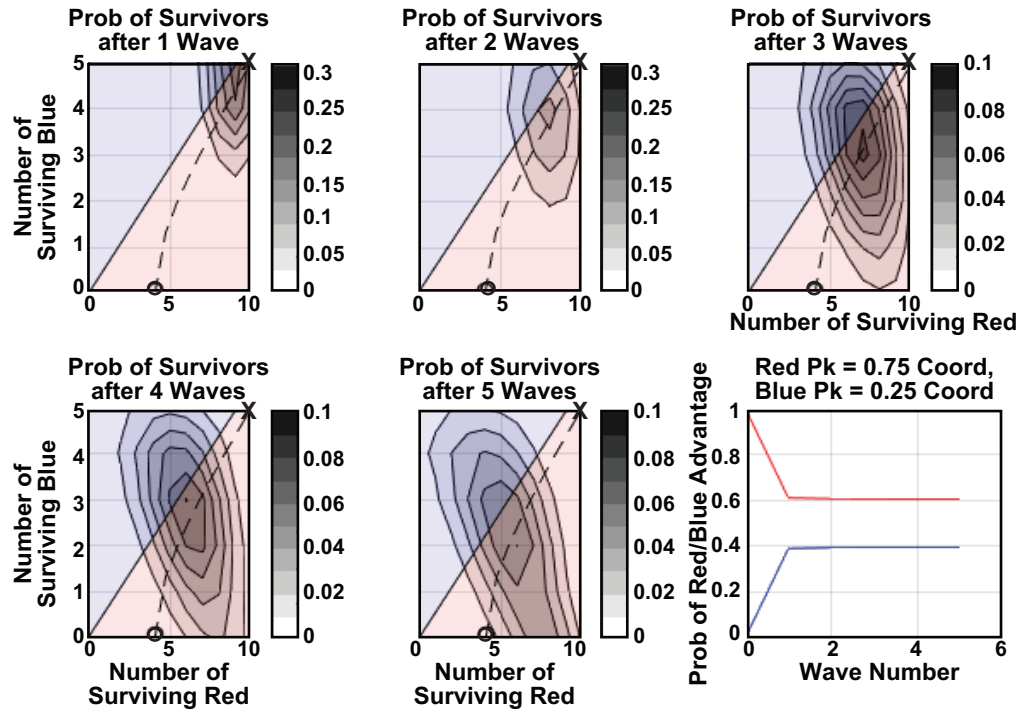


Figure 8. Results for reduced lethality.

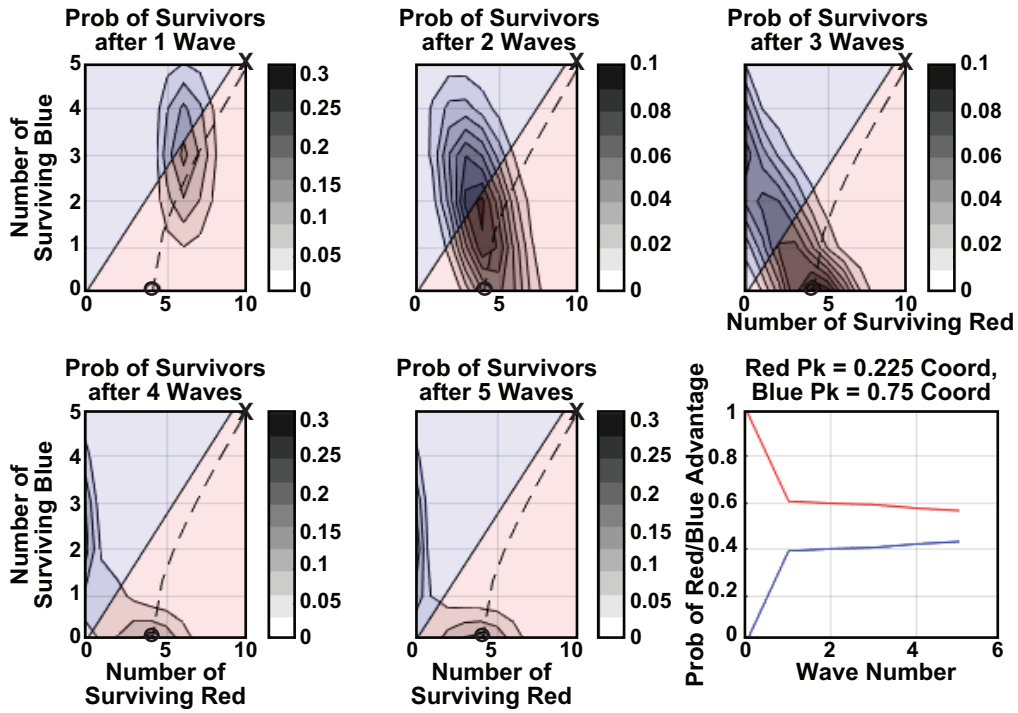


Figure 9. Results for increased lethality.

4.4 EFFECT OF ASSIGNMENT APPROACH

Here we look at the effect of changing the assignment approach from coordinated to random for red, blue, or both. In all cases, we will compare with Figure 5 where both sides use coordinated assignments. In Figure 10, both sides use random assignment. The results are quite similar to the case where both use coordinated assignment with a slight shift towards red winning.

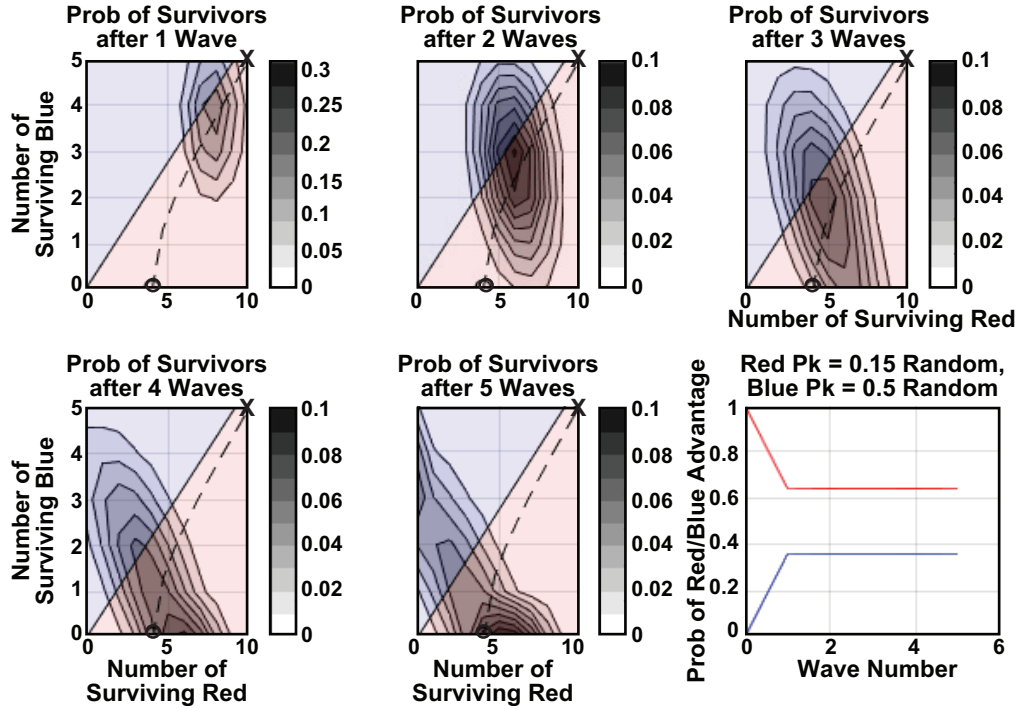


Figure 10. Both sides use random assignment.

In Figure 11, red uses a random assignment while blue uses a coordinated assignment. Blue does a little better but the results are still pretty much the same as when both use coordinated assignment. In Figure 12, we reverse things and have red use a coordinated assignment while blue uses a random assignment. In this case, blue does quite a bit worse than in the other cases. It seems that the effect of assignment approach is more important for the numerically smaller side that has to rely on its quality to make up for its lack of quantity. The numerically weaker side can't afford to waste any of its shots.

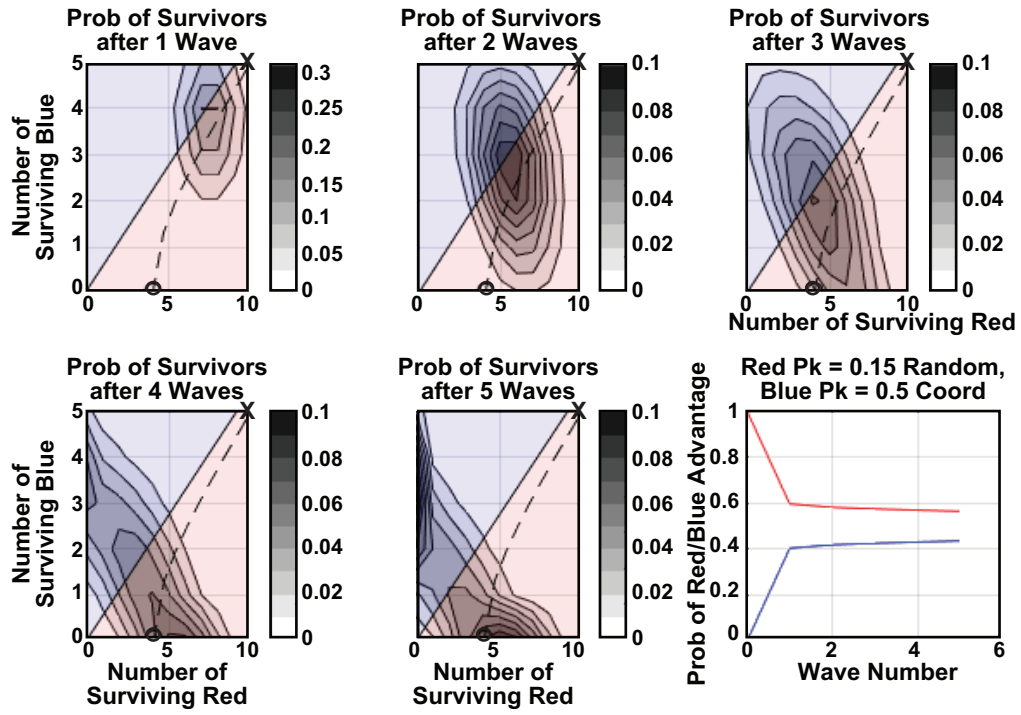


Figure 11. Red random, blue coordinated.

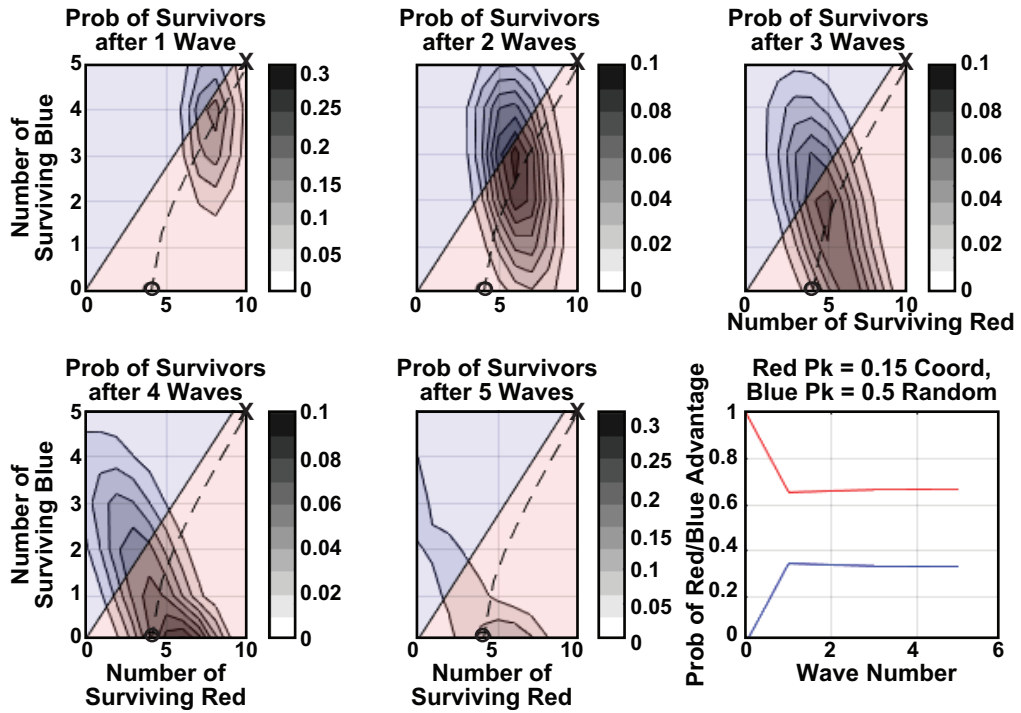


Figure 12. Red coordinated, blue random.

5. MODEL FOR RED AND BLUE WIN PROBABILITY

The previous plots have shown how the probability distribution of red and blue survivors evolves during the attack for different parametric variations. Next we will look at how the final probability of a red or blue win depends on the initial force levels for different parametric cases. We will make contour plots of the red win probability as a function of the initial red and blue force levels. We will have different contour plots for different combinations of assignment approaches and for different values of the quality ratio(P_{kr}/P_{kb}) of the opposing forces. Figure 13 shows the results for our nominal case shown in Figure 5 with $P_{kr} = 0.15$, $P_{kb} = 0.5$ and both sides using coordinated assignment.




Figure 13. Contours of red win probability.

The white line is the break even curve for the deterministic Lanchester model. We see a transition from a sure red win to a sure blue win over about a factor of 2 in the initial blue force level (or initial red force level). The next figures repeat Figure 13 for all combinations of weapon allocation strategies. Figure 14 is for the same kill probability ratios as Figure 13, and Figure 15 is for equal kill probabilities. In both figures, the difference between the different strategies is relatively small compared with the effects of randomness. The biggest difference is between the upper right and lower left plots where one side uses coordinated firing and the other uses random firing. In Figure 14, the red win probability on the break even line is close to 0.5 for red coordinated and blue random, but has decreased to about 0.4 for red random and blue coordinated. It looks like red suffers more for using the random strategy. In Figure 15 where the kill probabilities are equal, the red win probability at the break even point varies from 0.4 when red is random to 0.6 when blue is random. For this case, we would expect symmetry in the outcomes.

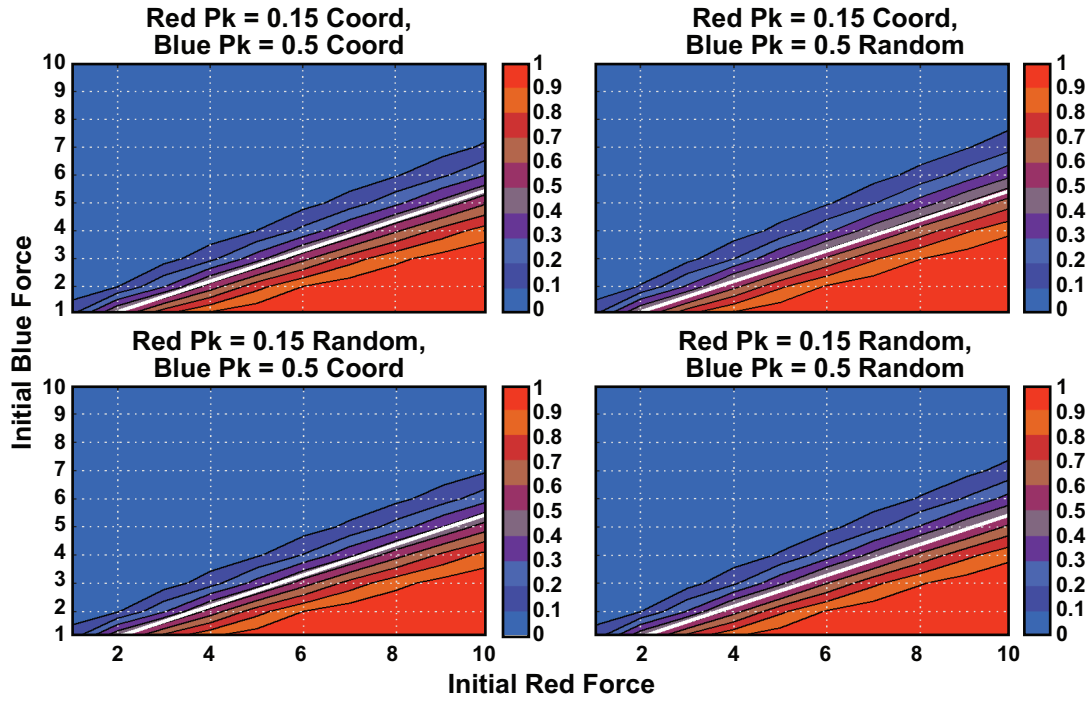


Figure 14. Red win probability for different weapon allocations.




Figure 15. Red win probability for equal lethalties.

5.1 MODELING THE PROBABILISTIC RESULTS

The fact that the contour plots in Figures 13–15 are relatively simple leads us to ask if we can come up with a simple model to estimate the red win probability as a function of the relative force levels and kill probabilities (and possibly the different weapon allocation strategies). In Figure 16, we plot the probability of a red win as a function of the initial force effectiveness ratio. We have taken hundreds of combinations of n_{red} , n_{blue} , P_{kr} , and P_{kb} , and calculated the effectiveness ratio $(n_{red}/n_{blue})^2(P_{kr}/P_{kb})$ for each case. We also calculated the probability of a red win for each case. The left hand plot is a scatter plot of these values where black dots represent cases with both sides using coordinated firing; green dots represent cases with both sides using random firing; and red and blue dots represent red or blue using coordinated firing and the other side using random firing. In the right hand plot, we convert the red win probability to a sigma value assuming a Gaussian distribution. For example, a red win probability of 0.5 corresponds to 0 sigma. A red win probability of 0.84 (or 0.16) corresponds to a 1 (or –1) sigma. If we plot the number of sigmas vs the log of the effectiveness ratio, we come close to a straight line. It is clear that there is a small effect due to the different firing doctrines, and it is in the direction that we would expect—coordinated firing is a little better than random firing. However, there is still significant scatter in the points about a best fit straight line.




Figure 16. Red win probability vs. effectiveness ratio.

Figure 17, shows the results for different ranges of threat levels. We see a clear difference in the results for small force levels and large force levels. There is less scatter about the best fit line when results are grouped by force level. These results are only for the case of both sides using coordinated firing. The computation requirements for this case are more relaxed than the requirements for one or both sides using random firing. This enables us to get results for much larger force levels. However, even these cases take a lot of time for larger force levels. The fact that the straight lines fit so well indicates that we can come up with a simple Gaussian approximation to the red win probability that will let us extend results to cases that would be too hard to calculate more rigorously.




Figure 17. Red win probability for different force levels.

Figure 18 shows how the approximate model works. The upper left plot shows how the slope of the lines in Figure 17 varies with the mean initial force levels. The upper right plot repeats the upper left plot on log-log paper. A simple relationship between the slopes and the force levels is evident. The lower left plot uses this simple model to regenerate the curves in Figure 17. Finally, the lower right plot converts the number of sigmas to the probability of a red win.




Figure 18. Approximate model.

The process used in the approximate model is outlined in Figure 19, starting with the force levels and effectiveness levels. All of these levels are used to generate a force effectiveness ratio.

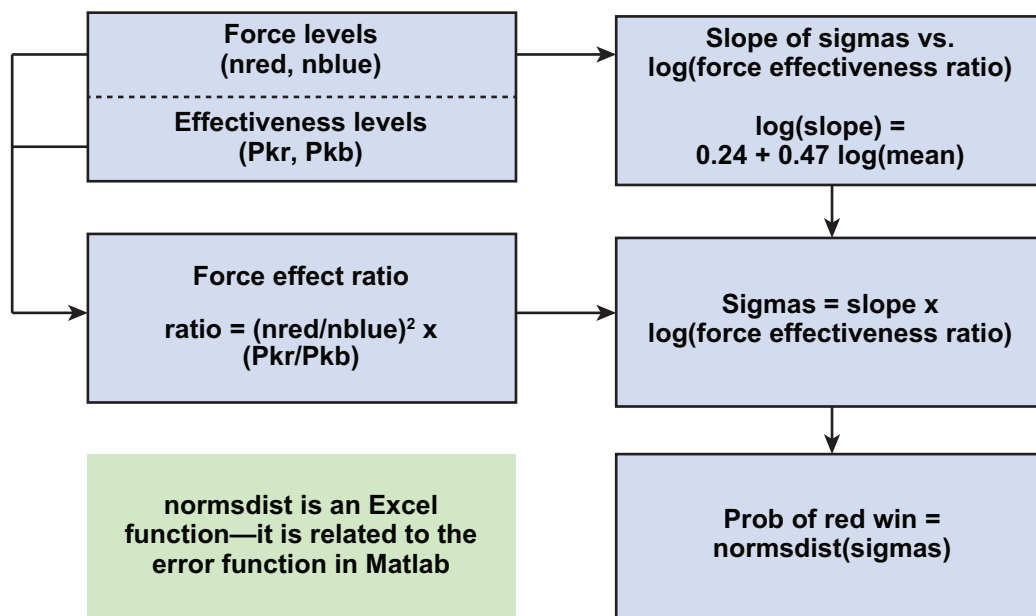


Figure 19. Overall approximate model.

In Figure 20, the results shown in Figure 13 are compared with results from the approximate model for the same parameter values. The results are very close, encouraging us to apply the approximate model to cases with much larger forces that would be impractical to calculate using the original model.




Figure 20. Comparison approximate model.

This is done in Figure 21 for force levels up to 100 on each side.




Figure 21. Approximate model for large force levels.

The blue force level is plotted absolutely on the left, and as a fraction of the red force level on the right. In the left plot, a greater range of blue force levels produces intermediate red win probabilities as the red force level increases. However, in the right plot, the relative range of blue force levels (relative to red) decreases as the red force level increases. This is a reflection of the increase in slope in Figure 17 as we increase the red force level. In the limit of very large forces, the relative range of blue to give intermediate win probabilities would shrink to be close to the deterministic case. This would occur only for much larger force levels than shown here.

This page intentionally left blank

6. SUMMARY

We have presented a probabilistic version of the classical Lanchester combat model. While the classical model is deterministic and the winner is determined only by the relative force levels and their relative qualities, the probabilistic model can result in a finite probability of the weaker side winning. This is particularly significant for modest force levels. As the force levels increase, the probabilistic results approach the deterministic results. We have also derived very simple analytic approximations to the probability of red or blue wins as functions of the force levels and qualities.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) TBD		2. REPORT TYPE Technical Report		3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE A Simple Probabilistic Combat Model				5a. CONTRACT NUMBER FA8721-05-C-0002 and/or FA8702-15-D-0001	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Stephen D. Weiner				5d. PROJECT NUMBER 9-033101	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) MIT Lincoln Laboratory 244 Wood Street Lexington, MA 02420-9108				8. PERFORMING ORGANIZATION REPORT NUMBER TR-1210	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force AFLCMC/PZM Bhupender Singh 20 Schilling Circle Bldg. 1305, 3 rd Fl. Hanscom AFB, MA 01731				10. SPONSOR/MONITOR'S ACRONYM(S) DHS	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT					
13. SUPPLEMENTARY NOTES					
<p>This report surveys different approaches to securely storing and sharing data in the cloud based on traditional notions of security: confidentiality, integrity and availability, with the main focus on confidentiality. An appendix discusses the related notion of how users can securely authenticate to cloud providers.</p> <p>We propose a metric for comparing secure storage approaches based on their <i>residual vulnerabilities</i>: attack surfaces against which an approach cannot protect. Our categorization therefore ranks approaches from the weakest (the most residual vulnerabilities) to the strongest (the fewest residual vulnerabilities). In addition to the security provided by each approach, we also consider their inherent costs and limitations. This report can therefore help an organization select a cloud data protection approach that satisfies their enterprise infrastructure, security specifications, and functionality requirements.</p>					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (include area code)